Let's Do Smart Routing: For Distributed Graph Querying with Decoupled Storage

Arijit Khan[1] Gustavo Segovia[2] Donald Kossmann[2,3]
[1]Nanyang Technological University, Singapore [2]ETH Zurich, Switzerland [3]Microsoft Research Redmond, USA

arXiv:1611.03959v1 [cs.DB] 12 Nov 2016

ABSTRACT

We study online graph queries that retrieve nearby nodes of a query node from a large network, e.g., find all 2-hop neighbors of Alice in the LinkedIn social network who are currently employed by Google. To answer such queries with high throughput and low latency, we partition the graph and process the data in parallel across a cluster of servers. State-of-the-art distributed graph querying systems place each graph partition on a separate server, where query answering over that partition takes place. This design has two major disadvantages. First, the router needs to maintain a fixed routing table. Hence, these systems are less flexible with respect to query routing, fault tolerance, and graph updates. Second, the graph data must be partitioned such that the workload across the servers is balanced, and the inter-machine communication is minimized. In addition, to maintain good-quality partitions, it is required to update the existing partitions based on workload changes over graph nodes. However, graph partitioning, online monitoring of workloads, and dynamically updating the graph partitions are expensive.

In this work, we mitigate both these problems of state-of-the-art systems by decoupling graph storage from query processors, and by developing smart routing strategies with the usage of graph landmarks and embedding, that improve the cache locality in query processors. Since a query processor is no longer assigned any fixed part of the graph, it is equally capable of handling any request, thus facilitating load balancing and fault tolerance. On the other hand, due to our smart routing strategies, query processors can effectively leverage their cache contents, reducing the overall impact of how the graph is partitioned across storage servers. A detailed experimental evaluation with several real-world, large graph datasets demonstrates that our proposed framework, gRouting — even with simple hash partitioning of the data — achieves up to an order of magnitude better query throughput compared to existing graph querying systems that employ expensive graph partitioning and re-partitioning strategies.

1. INTRODUCTION

Graphs with millions of nodes and billions of edges are ubiquitous to represent highly interconnected structures including the World Wide Web, social networks, knowledge graphs, genome and scientific databases, Internet of things, medical and government records. In order to support online search and query services (possibly from many clients) with low latency and high throughput, data centers and cloud operators consider scale-out solutions, in which the graph and its data are partitioned horizontally across cheap commodity servers in the cluster. We assume that the graph topology and the data associated with nodes and edges are co-located, since they are often accessed together [23, 24, 50]. Keeping with the modern database trends to support low-latency operations, we target a fully in-memory system, and use disks only for durability [37, 41, 53, 54].

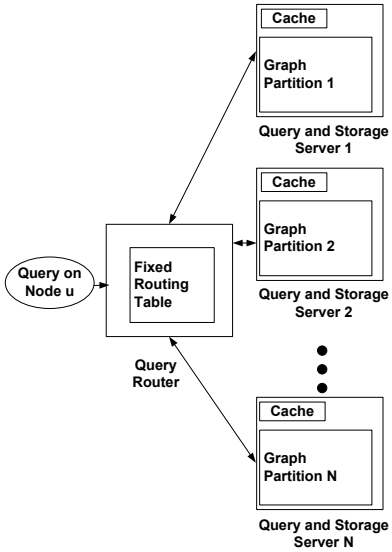

Figure 1: State-of-the-art distributed graph querying systems (e.g., SEDGE [53], Trinity [41], Horton [37])

In this paper, we study online queries that explore a small region of the entire graph, and require fast response time. These queries usually start with a query node, and traverse its neighboring nodes up to a certain number of hops. While we shall formally introduce our queries in Section 2, typical real-world examples can be as follows. **(1)** *Home Timeline Query* [50]: When a user logs in to Facebook or refreshes her homepage, a timeline is shown with recent posts created by her friends. **(2)** *Ego-Centric Query* [24]: In the LinkedIn social network, user Alice may search for her connections within 2-hops who are currently employed by Google. **(3)** *Pattern Matching Query* [23]: Given the Microsoft Academic graph, find all papers on "distributed graph systems" which are a result of collaboration between researchers from UC Berkeley and CMU. We generalize such online queries with an h-hop traversal from the query node(s).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2016 ACM. ISBN 978-1-4503-2138-9…$15.00
DOI: 10.1145/1235

For efficient answering of online h-hop queries in a distributed environment, state-of-the-art systems (e.g., [37, 41, 53]) first partition the graph, and then place each partition on a separate server, where query answering over that partition takes place (Figure 1). Since the server which contains the query node can only handle that request, the router always needs to maintain a fixed routing table (or, a fixed routing function, e.g., modulo hashing). Hence, these systems are less flexible with respect to query routing and fault tolerance, e.g., adding more machines will require updating the routing table. Besides, an effective graph partitioning in these systems must achieve: (1) workload balancing to maximize parallelism, and (2) locality of data access to minimize network communication. It has been shown in the literature [22, 53] that sophisticated partitioning schemes improve the performance of graph querying, compared to an inexpensive hash partitioning.

However, due to power-law degree distribution of real-world graphs, it is difficult to partition them well [9]. Besides, a one-time partitioning cannot cope with later updates to graph structure or variations in query workloads. Several graph re-partitioning and replication-based strategies were proposed to alleviate these problems, e.g., [22, 23, 28, 48, 53]. However, online monitoring of workload changes, re-partitioning of the graph topology, and migration of graph data across servers are expensive; and they reduce the efficiency and throughput of online querying [36].

Our Contribution. In contrast to existing systems, we consider a design with a different architecture, which relies *less* on an effective graph partitioning. Instead, we decouple query processing and graph storage into two separate tiers (Figure 2). In a decoupled framework, the servers allocated to the storage tier hold all the graph data in their main memory. Since a query processor is no longer assigned any fixed part of the graph, it is equally capable of handling any request, thus facilitating load balancing and fault tolerance. At the same time, the query router can send a request to any of the query processors, which adds more flexibility to query routing, e.g., more query processors can be added without affecting the routing strategy. Another benefit due to decoupled design is that each tier can be scaled-up independently. If a certain workload is processing intensive, more servers could be allocated to the processing tier. On the other hand, if the graph size increases over time, more servers can be added in the storage tier. This decoupled architecture, being generic, can be employed in many existing graph querying systems.

The idea of decoupling, though effective, is not novel. Facebook implemented a fast caching layer, Memcached on top of a graph database that scales the performance of graph query answering [29]. Google's F1 [43] and ScaleDB (http://scaledb.com/pdfs/TechnicalOverview.pdf) are based on a similar decoupling principle for scalability. Very recently, Loesing et. al. [19] and Binnig et. al. [3] demonstrated the benefits of a decoupled, shared-data architecture, together with low latency and high throughput Infiniband network. Shalita et. al. [39] employed de-coupling for an optimal assignment of HTTP requests over a distributed graph storage.

Our fundamental contribution lies in designing a smart query routing logic to utilize the cache of query processors over such decoupled architecture. Achieving more cache hits is critical in a decoupled architecture – otherwise, the query processors need to retrieve the data from storage servers, which will incur extra communication costs. However, this is a non-trivial problem, e.g., exploiting cache locality and balancing workloads are conflicting in nature. In addition, graph querying workloads are significantly different from traditional database applications. The interconnected nature of graph data results in poor locality, and each query usually accesses multiple neighboring nodes spreading across the distributed storage. Therefore, to maximize cache hit rates at query processors, it is not sufficient to only route the queries on same nodes to the same processor. Rather, successive queries on *neighboring* nodes should also be routed to the same processor, since the neighborhoods of two nearby nodes may significantly overlap. Besides, the smart routing logic that we design to resolve the above challenges, should be lightweight on storage and online computation cost, as well as adaptive with workload hostspot and graph updates. To the best of our knowledge, such smart query routing schemes for effectively leveraging the cache contents were not considered in state-of-the-art graph querying systems, including SEDGE [53], Trinity [41], Horton [37]. *In this paper, we take the first step to this goal*.

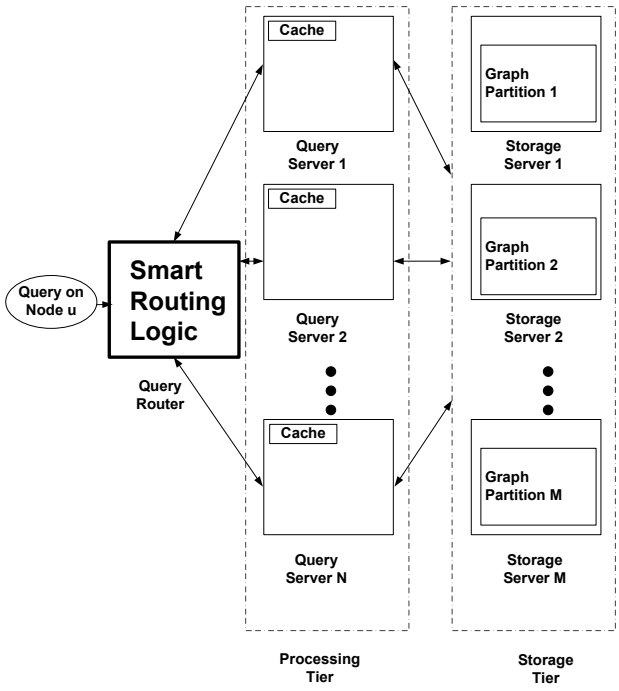

Figure 2: Decoupled architecture for distributed graph querying

We summarize our contributions as follows.

1. We investigate for the first time the novel problem of query routing aimed at improving the throughput and efficiency of distributed graph querying.
2. In contrast to many recent distributed graph querying systems [37, 41, 53], we consider a generic architecture that decouples query processors from storage layer, thereby achieving flexibility in system deployment, query routing, scaling up, load balancing, and fault tolerance.
3. We develop smart, lightweight, and adaptive query routing algorithms that improve cache hit rates at query processors, thereby reducing communication with storage layer, and making our design less reliant on a sophisticated graph partitioning scheme across storage servers.
4. We empirically demonstrate throughput and efficiency of our framework, gRouting on four real-life graphs, while also comparing with two state-of-the-art distributed graph processing systems (SEDGE/Giraph [53] and PowerGraph [9]). Our decoupled implementation, even with its simple hash partitioning, achieves up to an order of magnitude higher throughput compared to these existing systems with expensive graph partitioning and re-partitioning schemes.

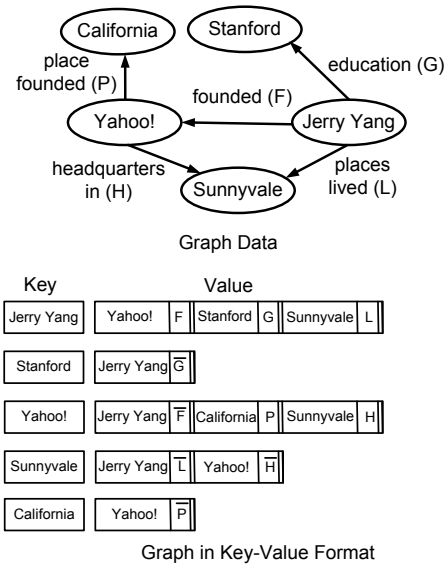

Figure 3: Key-Value storage format of graph data: $\bar{F}$ denotes the inverse relationship of F.

2. PRELIMINARIES

2.1 Graph Data Model

A heterogeneous network can be modeled as a labeled, directed graph $G = (V, E, \mathcal{L})$ with node set V, edge set E, and label set L, where (1) each node $u \in V$ represents an entity in the network, (2) each directed edge $e \in E$ denotes the relationship between two entities, and (3) $\mathcal{L}$ is a function which assigns for each node u and every edge e a label $\mathcal{L}(u)$ (and $\mathcal{L}(e)$, respectively) from a finite alphabet. In practice the node labels may represent the attributes of the entities, such as name, job, location, etc, and edge labels the type of relationships, e.g., founder, place founded, etc.

We store the graph data as an adjacency list (Figure 3). Every node in the graph is added as an entry in the storage where the key is the node id and the value is an array of 1-hop neighbors. If the nodes and edges have labels, they are also stored in the corresponding value entry. For each node, we store both incoming and outgoing edges. This is because both incoming and outgoing edges of a node can be important from the context of different queries [1]. As an example, if there is an edge *founded* from *Jerry Yang* to *Yahoo!* in some knowledge graph, there also exists a reverse relation *founded_by* from *Yahoo!* to *Jerry Yang*. Such information could be useful in answering queries about *Yahoo!*.

2.2 h-Hop Traversal Queries

We introduce various h-hop queries [2] over heterogeneous, directed graphs in the following.

1. h-**hop Neighbor Aggregation**: Count the number of h-hop neighbors (or, the number of occurrences of a specific label within the h-hop neighborhood) of a query node.
2. h-**step Random Walk with Restart**: The query starts at a node, and at each successive step jumps to one of its neighbors with equal probability, or returns to the query node with a small probability. The query consists of h steps.
3. h-**hop Reachability**: Given a source and a target node, find if the target is reachable from the source within h-hops.

[1] Since we store both incoming and outgoing edges for every node, we shall consider every edge as a bi-directed edge in our routing algorithms in Section 3.4.

The aforementioned queries are often used as the basis for more complex graph operations. For example, neighborhood aggregation is critical for node labeling and classification [38], that is, the label of an unlabeled node could be assigned as the most frequent label which is present within its h-hop neighborhood. The h-step random walk is useful in expert finding, ranking, discovering functional modules, complexes, and pathways, and also in personalized PageRank computation [12, 20]. For processing the third query, we perform a bi-directional breadth first search (BFS), i.e., in forward direction from the source node and in backward direction from the target node, which is feasible because we store both incoming and outgoing edges for every node in the graph dataset. In addition, if there are node- and edge-label constraints in reachability computation, e.g., [13], one can enforce that while performing the BFS. Therefore, our third query can be employed in distance-constrained and label-constrained reachability search, as well as in approximate graph pattern matching queries [14, 45].

2.3 Decoupled Design

We review the general implementation details of our decoupled architecture — discussing their merits, and how to overcome their limitations, which will lead to introducing our smart query routing algorithms in the following section.

We decouple query processing from graph storage. One may note that this decoupling happens at a *logical* level. As an example, query processors can be different physical machines than storage servers. On the other hand, the same physical machine can also run a query processing daemon, together with storing a graph partition in its main memory as a storage server. However, the logical separation between the two layers is critical in our design, i.e., query processors do not keep a partition of the graph, and must access data storage to retrieve any part of the graph that they need for answering a query.

As stated earlier, the advantages of this separation are more flexibility in query routing, system deployment, and scaling up, as well as achieving better load balancing and fault tolerance. However, we must also consider the drawbacks of having the graph storage apart. First, query processors may need to communicate with the storage tier via the network. This includes an additional penalty to the response time for answering a query. Second, it is possible that this design causes high contention rates on either the network, storage tier, or both.

To mitigate these issues, we design smart routing strategies that route online queries to processors which are likely to have the relevant data in their cache, thereby reducing the communication overhead between processing and storage tiers. Below, we discuss various components of our design, including storage, processing tier, and router.

Graph Storage Tier. The storage tier is responsible for holding all the graph data by horizontally partitioning it across cheap commodity servers. A sophisticated graph partitioning scheme could benefit our decoupled architecture in the following way. Let us assume that the neighboring nodes can be stored in a page within the same storage server, and the granularity of transfer from storage to processing tier is a page containing several nodes. Then, we could actually ship a set of relevant nodes with a single request if the graph is partitioned well. This will reduce the number of times data are transferred between the processing and storage tier.

However, our lightweight and smart query routing techniques exploit the notion of graph landmarks [17] and embedding [56], thereby effectively utilizing the cache of query processors that

[2] An algorithm to evaluate h-hop BFS queries in gRouting is given in Appendix.

stores recently used graph data. As demonstrated in our experiments, most part of a query can be computed locally in the query processors using their cache contents. Therefore, the partitioning scheme employed across storage servers becomes less important in our framework.

Query Processing Tier. The query processing tier consists of servers where the actual query processing takes place. These servers do not communicate with each other [19]. They only receive queries from the query router. They can also request graph data from the storage tier, if necessary.

To reduce the amount of calls made to the storage tier, we utilize the cache of the query processors. Whenever some data is retrieved from the storage, it is saved in cache, so that the same request can be avoided in the near future. However, it imposes a constraint on the maximum storage capacity. When the addition of a new entry surpasses this storage limit, one or more old entries are evicted from the cache. We chose the LRU (i.e., Least Recently Used) eviction policy because of its simplicity. LRU is usually implemented as the default cache replacement policy, and it favors recent queries over old ones. Thus, it performs well with our smart routing schemes.

Query Router. The router creates a thread for each processor and opens a connection to send queries by following some routing schemes that we describe in Section 3.

3. QUERY ROUTING STRATEGIES

When a query arrives at the router, it decides the appropriate query processor to which the request could be sent. For state-of-the-art graph querying systems, e.g., SEDGE [53] and Horton [37], where each query processor is assigned a graph partition, this decision is fixed and defined in the routing table; the query processor which contains the query node should handle the request. With a decoupled architecture, no such mapping exists, and hence, we design novel routing schemes with the following objectives.

3.1 Routing Algorithm Objectives

We aim at improving the query throughput and reducing the latency. To this end, we identify the following criteria that a *smart routing* scheme must possess.

1. Leverage each processor's cached data. To formalize this notion, let us consider t queries $q_1, q_2, \ldots, q_t$ received successively by the query router. The router will send them to query processors in a way such that the average number of *cache hits* at the processors is maximized. This, in turn, reduces the average query processing latency. However, as stated earlier, to achieve maximum cache hits, it will not be sufficient to only route the queries on same nodes to the same processor. Rather, successive queries on *neighboring* nodes should also be routed to the same processor, since the neighborhoods of two nearby nodes may significantly overlap. This will be discussed shortly in Requirement 1.

2. Balance workload even when it is skewed or contains hotspot. As earlier, let us consider a set of t successive queries. A naïve approach will be to ensure that each query processor receives equal number of queries, e.g., a round-robin way of query dispatching by the router. However, each query might have a different workload, and would require a different processing time. We, therefore, aim at maximizing the overall *throughput* via query stealing (explained in Requirement 2), which automatically balances the workload across query processors.

3. Make fast routing decisions. The average time at the router to dispatch a query should be minimized, ideally a small constant time, or much smaller than $\mathcal{O}(n)$, where n is the number of nodes in the input graph. This reduces the query processing latency.

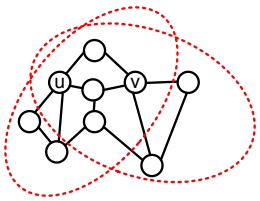

Figure 4: Topology-aware locality: 2-hop neighborhoods of nearby nodes u and v significantly overlap.

4. Have low storage overhead in the router. The router may store some auxiliary data structures to enable fast routing decisions. However, this additional storage overhead must be a small fraction compared to the input graph size.

3.2 Challenges in Query Routing

It is important to note that our routing objectives are not in harmony; in fact, they are often conflicting with each other. First, in order to achieve maximum cache locality, the router can send all the queries to the same processor (assuming no cache eviction happens). However, the workload of the processors will be highly imbalanced in this case, resulting in lower throughput. Second, the router could inspect the cache of each processor before making a good routing decision, but this will add network communication delay. Hence, the router must *infer* what is likely to be in each processor's cache.

In the following, we introduce two concepts that are directly related to our routing objectives, and will be useful in designing smart routing algorithms.

Topology-Aware Locality. To understand the notion of cache locality for graph queries (i.e., routing objective 1), we define a concept called *topology-aware locality*. If u and v are nearby nodes, then successive queries on u and v must be sent to the same processor. It is very likely that the h-hop neighborhoods of u and v will significantly overlap (Figure 4). We refer to this phenomenon as the topology-aware locality.

But, how will the router know that u and v are nearby nodes? One option is to store the entire graph topology in the router; but this could have a high storage overhead. For example, the *WebGraph* dataset that we experimented with has a topology of size 60GB. Ideally, a graph with 10^7 nodes can have up to 10^{14} edges, and in such cases, storing only the graph topology itself would require petabytes of memory. Thus, we impose a requirement for our smart routing schemes as follows.

REQUIREMENT 1. *The additional storage at the router for enabling smart routing should not be asymptotically larger than $\mathcal{O}(n)$, n being the number of nodes; however, the routing schemes should still be able to exploit topology-aware locality.*

One may realize that achieving this goal is non-trivial, since the topology information can be $\mathcal{O}(n^2)$, and we provision for only $\mathcal{O}(n)$ space to approximately preserve such information.

Query Stealing. We found that routing queries to processors that have the most useful cache data might not always be the best strategy. Due to power-law degree distribution of real-world graphs, processing queries on different nodes might require different amount of time. Therefore, the processors dealing with high-degree nodes will have more workloads. Load imbalance can also happen if the queries are concentrated in one specific region of the graph. When that happens, all queries will be sent to one processor, while other processors remain idle. To rectify such scenarios, we implement *query stealing* in all our routing schemes as stated next.

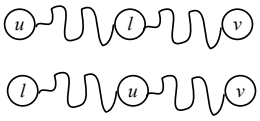

Figure 5: Strict upper (above) and lower (below) bounds on shortest-path distances based on landmark

REQUIREMENT 2. *Whenever a processor is idle and is ready to handle a new query, if it does not have any other requests assigned to it, it may "steal" a request that was originally intended for another processor.*

Query stealing is a well established technique for load balancing that is prevalently used by the HPC community, and there are several ways how one can implement it. The idle processors can actively steal queries from other processors that are currently busy [7]. In our implementation, we assume that the processors send acknowledgements to the router when they finish computation. Thus, *we perform query stealing at the router level*. In particular, the router sends the next query to a processor only when it receives an acknowledgement for the previous query from that processor. The router also keeps a *queue* for each connection in order to store the future queries that need to be delivered to the corresponding processor. By monitoring the length of these queues, it can estimate how busy a processor is, and this enables the router to rearrange the future queries for load balancing.

We next design four different routing schemes — the first two are naïve and do not meet all the objectives of smart routing. On the other hand, the last two algorithms follow the requirements of a smart routing strategy. We, therefore, use the first two routing schemes as baselines to compare against our smart routing algorithms in the experiments (Section 4).

3.3 Baseline Methods

3.3.1 Next Ready Routing

Next Ready routing is our first baseline strategy. The router decides where to send a query by choosing the next processor that has finished computing and is ready for a new request. The main advantages of this scheme are:

- It is easy to implement.
- Routing decisions are made in constant time.
- No preprocessing or storage overhead is required.
- The workload is well balanced.

Even though this scheme has several benefits, it fails to leverage processors' cache.

3.3.2 Hash Routing

The second routing scheme that we implement is *hash*, and it also serves as a baseline to compare against our smart routing techniques. The router applies a fixed hash function on each query node's id to determine the processor where it sends the request. In our implementation, we apply a modulo hash function:

$$\text{Target-Processor-Id} := (\text{Query-Node-Id}) \mathsf{MOD} (\text{Number-Of-Processors}) \quad (1)$$

In order to facilitate load balancing in the presence of workload skew, we implement *query stealing* mechanism. Whenever a processor is idle and is ready to handle a new query, if it does not have any other requests assigned to it, it steals a request that was originally intended for another processor. Since queries are queued in the router, the router is able to take this decision, and ensures that there are no idle processors when there is still some work to be done. Our hash routing has all the benefits of next ready, and very likely it sends a repeated query to the same processor, thereby getting better locality out of the cache. However, hash routing cannot capture topology-aware locality.

Algorithm 1 Preprocessing for Landmark Routing

Require: Bi-directed graph G, #landmarks = $|L|$, #query-processors = P
Ensure: Landmarks set L, distance $d(u, p)$ from every node u to each processor p
1: Find $|L|$ max degree nodes in G
2: **for all** landmark in L **do**
3: Perform BFS to find distance to each node in G
4: **if** distance between two landmarks $<$ threshold **then**
5: discard one with min degree
6: **end if**
7: **end for**
/* Find P pivot landmarks */
8: Find two landmarks farthest apart and mark them as pivots
9: **while** P pivots not found **do**
10: Find the landmark farthest from previous pivots, and mark it a pivot
11: **end while**
/* Assign landmarks to processors */
12: Assign every pivot landmark to a distinct processor
13: Assign remaining landmarks to processors having the closest pivot
/* Compute distance from nodes to processors */
14: **for all** node u in G and processor p **do**
15: $d(u, p) \rightarrow$ min distance between u to any landmark assigned to p
16: **end for**

3.4 Proposed Methods

3.4.1 Landmark Routing

Our first smart routing scheme is *Landmark*, based on the notion of landmark nodes [17, 32]. One may recall that we store both incoming and outgoing edges of every node, thus we consider a *bi-directed* version of the input graph in our smart routing algorithms. In particular, we assume a bi-directed edge corresponding to every directed edge in the input graph, then we select a small set L of nodes as landmarks, and also pre-compute the distance of every node to these landmarks. We determine the optimal number of landmarks based on empirical results. As depicted in Figures 5, given some landmark node $l \in L$, the distance $d(u, v)$ between any two nodes u and v are bounded as follows:

$$\begin{aligned} d(u, v) &\leq d(u, l) + d(l, v) \\ d(u, v) &\geq |d(u, l) - d(l, v)| \end{aligned} \quad (2)$$

Intuitively, if two nodes are close to a given landmark, they are likely to be close themselves. Our landmark routing is based on the above principle. We first select a set of landmarks that partitions the graph into P regions, where P is the total number of processors. We then decide a fixed one-to-one mapping between those regions and processors. Now, if a query belongs to a specific region (decided based on its distance to landmarks), it is routed to the corresponding processor. Clearly, this routing strategy requires a preprocessing phase (Algorithm 1), where we select the landmark nodes, compute the distance of every node to these landmarks, as well as decide the mapping between landmarks and processors.

Preprocessing. By following [32], we select landmarks based on their node degree and how well they spread over the entire graph. Our first step is to find a certain number of landmarks considering the highest degree nodes (line 1, Algorithm 1), and then compute their distance to every node in the graph by performing breadth first searches (BFS). In this process, if we find two landmarks to be

Algorithm 2 Landmark Routing Algorithm

Require: Query on node u
Ensure: Processor p^* to route query on u
1: **for all** processor p **do**
2: Compute $d^{LB}(u,p)$ following Equation 3
3: **end for**
4: $p^* \rightarrow \arg\min_p \left(d^{LB}(u,p)\right)$

closer than a pre-defined threshold, the one with the lower degree is discarded (lines 4-5, Algorithm 1). The complexity of this step is $\mathcal{O}(|L|e)$, due to $|L|$ number of BFS, where $|L|$ is the number of landmarks, and e is the number of edges.

Next, we assign the landmarks to query processors (lines 8-13, Algorithm 1) as follows. First, every processor is assigned a "pivot" landmark with the intent that these pivot landmarks are as far from each other as possible.

- The first two pivot landmarks are the two landmarks that are farthest apart considering all other landmark pairs.
- Each subsequent pivot landmark is selected as the one that is farthest from all previously selected pivot landmarks.

Each of the remaining landmarks is assigned to the processor which contains its closest pivot landmark. The complexity of this step is $\mathcal{O}(|L|^2 + |L|P)$, where P is the number of query processors.

Our final step of preprocessing is to define a "distance" metric d between the graph nodes and query processors (lines 14-15, Algorithm 1). The distance of a node u to a processor p is defined as the minimum distance of u to any landmark that is assigned to processor p. Thus, we compute a distance value $d(u,p)$ for every node u to every processor p. This information is stored in the router, which requires $\mathcal{O}(nP)$ space and $\mathcal{O}(nL)$ time to compute, where n is the number of nodes. Therefore, *the storage requirement in the router is linear to the number of nodes*.

Routing. To decide where to send a query on node u, the router verifies the pre-computed distance $d(u,p)$ for every processor p, and selects the one with the smallest $d(u,p)$ value. As a consequence, the routing decision time is linear to the number of processors: $\mathcal{O}(P)$. This is very efficient since the number of processors is small.

In contrast to our earlier baseline routings, this method is able to leverage topology-aware locality. It is likely that query nodes that are in the neighborhood of each other will have similar distances to the processors; hence, they will be routed in a similar fashion. On the other hand, the landmark routing scheme is less flexible with respect to addition or removal of processors, since the assignment of landmarks to processors, as well as the distances $d(u,p)$ for every node u and each processor p needs to be recomputed.

We note that the distance metric $d(u,p)$ is useful not only for finding the best processor for a certain query, but it can also be used for load balancing, fault tolerance, dealing with workload skew, and hotspots. For example, let us assume that the closest processor for a certain query is very busy, or is currently down. Since the distance metric gives us distances to all processors, the router is able to select the second, third, or so on closest processor. This form of load balancing will impact the nearby query nodes in the same way; and therefore, the modified routing scheme will still be able to capture topology-aware locality. However, in practice, it can be complex to define exactly when a query should be routed to its next best query processor. We propose the following formula that calculates the *load-balanced distance* $d^{LB}(u,p)$ between a query node u and a processor p.

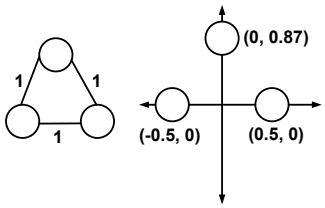

Figure 6: Example of graph embedding in 2D Euclidean plane

$$d^{LB}(u,p) := d(u,p) + \frac{\text{Processor Load}}{\text{Load Factor}} \tag{3}$$

Thus, the query is always routed to the processor with the smallest $d^{LB}(u,p)$ value, as presented in Algorithm 2. The router uses the number of queries in the queue corresponding to a processor as the measure of its load. The load factor is a tunable parameter, which allows us to decide how much load would result in the query to be routed to another processor. We find the optimal value of load factor via experiments.

Dealing with Graph Updates. In case of graph updates due to addition/ deletion of nodes and edges, one needs to recompute the distances from every node to each of the landmarks. This can be performed efficiently by keeping an additional *shortest-path-tree* data structure [46]. However, in order to avoid the additional space and time complexity of maintaining a shortest-path-tree, we follow a simpler approach. When a new node u is added, we compute the distance of this node to every landmark, and also its distance $d(u,p)$ to every processor p. In case of an edge addition or deletion between two existing nodes, for these two end-nodes and their neighbors up to a certain number of hops (e.g., 2-hops), we recompute their distances to every landmark, as well as to every processor. Finally, in case of a node deletion, we handle it by considering deletion of multiple edges that are incident on it. After a significant number of updates, previously selected landmark nodes become less effective; thus, we recompute the entire preprocessing step periodically in an off-line manner.

3.4.2 *Embed Routing*

Our second smart routing scheme is called the *Embed* routing, which is based on a graph embedding principle [6,56]. We embed a graph into a lower dimensional Euclidean space such that the hop-count distance between graph nodes are approximately preserved via their Euclidean distance (Figure 6). As depicted in [56] and also found in our experiments, higher dimensions produce smaller errors in terms of preserving node-pair distances. Particularly, we observed that an embedding of dimensionality 10 or higher is able to preserve the distances reasonably well. We then use the resulting node co-ordinates to determine how far a query node is from the recent history of queries that were sent to a specific processor. Clearly, embed routing also requires a preprocessing step as discussed in the following.

Preprocessing. We outline our preprocessing technique in Algorithm 3. For efficiently embedding a large graph in a D-dimensional Euclidean plane, we first select a set L of landmarks and find their distances from each node in the graph (lines 1-4, Algorithm 3). We then assign co-ordinates to landmark nodes such that the distance between each pair of landmarks is approximately preserved. We, in fact, minimize the *relative error* in distance for each pair of landmarks (line 5, Algorithm 3), defined below.

$$f_{error}(v_1,v_2) := \frac{|d(v_1,v_2) - \text{EuclideanDist}(v_1,v_2)|}{d(v_1,v_2)} \tag{4}$$

Algorithm 3 Preprocessing for Embed Routing

Require: Bi-directed graph G, #landmarks = $|L|$, #dimension = D
Ensure: Co-ordinates of every node u in G
1: Find $|L|$ max degree nodes in G
2: **for all** landmark in L **do**
3: Perform BFS to find distance to each node in G
4: **end for**
5: Embed landmarks: min rel. dist. error for all landmark pairs
6: **for all** non-landmark node u **do**
7: Embed u: min rel. dist. error between u and all landmarks
8: **end for**

Algorithm 4 Embed Routing Algorithm

Require: Query on node u, MeanCo-ordinate(p) for all processors p
Ensure: Processor p^* to route query on u
1: **for all** processor p **do**
2: Compute $d_1^{LB}(u,p)$ following Equation 7
3: **end for**
4: $p^* \rightarrow \arg\min_p \left(d_1^{LB}(u,p)\right)$
5: Update MeanCo-ordinate(p^*) following Equation 5

Here, $d(v_1, v_2)$ is the hop-count distance between v_1 and v_2 in the original graph, and EuclideanDist(v_1, v_2) is their Euclidean distance after the graph is embedded. We minimize the relative error since we are more interested in preserving the distances between nearby node pairs. Our problem is to minimize the aggregate of such errors over all landmark pairs — this can be cast as a generic multi-dimensional global minimization problem, and could be approximately solved by many off-the-shelf techniques, e.g., the *Simplex Downhill* algorithm that we apply in this work. Next, every other node's co-ordinates are found also by applying the Simplex Downhill algorithm that minimizes the aggregate relative distance error between the node and all the landmarks (lines 6-8, Algorithm 3). The overall graph embedding procedure consumes a modest preprocessing time: $\mathcal{O}(|L|e)$ due to BFS from $|L|$ landmarks, $\mathcal{O}(|L|^2 D)$ for embedding the landmarks, and $\mathcal{O}(n|L|D)$ for embedding the remaining nodes. In addition, the second step is completely parallelizable per node. Since each node receives D co-ordinates, it requires total $\mathcal{O}(nD)$ space in the router, which is linear to the number of nodes.

Unlike landmark routing, *a benefit of embed routing is that the preprocessing is independent of the system topology, allowing, for example, more processors to be easily added at a later time*.

Routing. The router has access to each node's co-ordinates. By keeping an average of the query nodes' co-ordinates that it sent to each processor, the router would be able to infer the cache contents in these processors. As a consequence, the router can find the distance between a query node u and a processor p, denoted as $d_1(u,p)$, and defined as the distance of the query node's co-ordinates to the historical mean of the processor's cache contents. Since recent queries are more likely to influence the cache contents due to our LRU eviction policy, we use the exponential moving average (EMA) function to compute the mean of the processor's cache contents [25]. Initially, the mean co-ordinates for each processor are assigned uniformly at random. Next, assuming that the last query on node v was sent to processor p, its mean co-ordinates are updated as follows.

$$\begin{aligned}&\text{MeanCo-ordinates}(p)\\ &:= \alpha \cdot \text{MeanCo-ordinates}(p) + (1-\alpha)\cdot \text{Co-ordinates}(v)\end{aligned} \tag{5}$$

The smoothing parameter $\alpha \in (0,1)$ in the above Equation determines the degree of decay used to discard older queries. For example, α close to 0 assigns more weight only to the last query, and α close to 1 decreases the weight on the last query. We determine the optimal value of α based on experimental results. Finally, the distance between a query node u and a processor p is computed as given below.

$$d_1(u,p) := ||\text{MeanCo-ordinates}(p) - \text{Co-ordinates}(u)|| \tag{6}$$

Since we consider embedding in an Euclidean plane, we use the L_2 norm in the above equation to compute distances. We select the processor with the smallest $d_1(u,p)$ distance. One may observe that the routing decision time is only $\mathcal{O}(PD)$, P being the number of processors and D the number of dimensions.

Analogous to landmark routing, we now have a distance to each processor for a query; and hence, we are able to make routing decisions taking into account the processors' workloads and faults. As earlier, we define a *load-balanced distance* $d_1^{LB}(u,p)$ between a query node u and a processor p, and the query is always routed to the processor with the smallest $d_1^{LB}(u,p)$ value.

$$d_1^{LB}(u,p) := d_1(u,p) + \frac{\text{Processor Load}}{\text{Load Factor}} \tag{7}$$

The embed routing, depicted in Algorithm 4, has all the benefits of smart routing. In fact, this routing strategy divides the active regions (based on query workloads) of the graph into P partitions in an overlapping manner, and assigns them to the the processors' cache. Moreover, it dynamically adapts the partitions with new workloads. Therefore, *with this smart embed routing, we are able to bypass the expensive graph partitioning and re-partitioning problems to the existing cache replacement policy of the query processors*. This shows the effectiveness of our embed routing.

Dealing with Graph Updates. Due to pre-assignment of node co-ordinates, embed routing is less flexible with respect to graph updates. When a new node is added, we compute its distance from the landmarks, and then assign co-ordinates to the node by applying the Simplex Downhill algorithm. Edge updates and node deletions are handled in a similar method as discussed for landmark routing. We recompute the entire preprocessing step periodically in an off-line manner to deal with a significant number of graph updates.

4. EXPERIMENTAL RESULTS

We present experimental results with four real-world large graphs. We evaluate online query processing efficiency, throughput, and cache usage of gRouting under various query workloads. We analyze the deployment flexibility and scalability by varying the number of query processors and storage servers. We also measure adaptiveness of our smart routing techniques with dynamic graph updates. We provide sensitivity analysis of our routing algorithms by varying several important parameters such as load factor, number of landmarks and their distances, embedding dimension, and smoothing parameter in the Appendix.

4.1 Compared Systems

To the best of our knowledge, smart query routing logic for distributed graph querying is novel. Thus, no direct competitors exists in the literature. Nevertheless, we compare gRouting with two state-of-the-art distributed graph processing systems: SEDGE/Giraph [53] and PowerGraph [9]. Other recent graph querying systems, e.g., [29, 37] are not publicly available for a direct comparison. However, decoupled architecture and our smart routing logic, being generic, can benefit many of them.

SEDGE/Giraph. SEDGE [53] was developed to support h-hop traversal queries (i.e., random walk, reachability, and neighborhood

Dataset	# Nodes	# Edges	Size on Disk (Adj. List File)
WebGraph	105 896 555	3 738 733 648	60.3 GB
Friendster	65 608 366	1 806 067 135	33.5 GB
Memetracker	96 608 034	418 237 269	8.2 GB
Freebase	49 731 389	46 708 421	1.3 GB

Table 1: Graph datasets

aggregation), on top of Giraph or Google's Pregel system [21], which follows in-memory, vertex-centric, bulk-synchronous parallel model. The authors of SEDGE employed ParMETIS software [15] for graph partitioning and re-partioning. While graph partitioning is generic and is not optimized for any specific workload, their re-partitioning method requires prior knowledge about the future queries. Hence, re-partitioning is a more sophisticated graph partitioning that is optimized for future workloads. One may note that our gRouting framework is adaptive with workload changes, and it does not require any prior knowledge of future queries for smart query routing. Nevertheless, we compare against SEDGE with re-partitioning, since SEDGE provides the best throughput with re-partitioned graphs.

PowerGraph. PowerGraph [9] follows in-memory, vertex-centric, asynchronous gather-apply-scatter model. For efficient implementation of our h-hop traversal queries, we ensure that *only* the required nodes are active at any point of time, that is, in the beginning, only the query node is active, and each active node then activates its neighbors, until all the h-hop neighbors from the query nodes are activated. PowerGraph also employs a sophisticated node-cut based graph partitioning method, for improved query throughput.

4.2 Environment Setup

• **Cluster Configuration.** We perform experiments on a cluster of 12 servers having 2.4 GHz Intel Xeon processors, and interconnected by 40 Gbps Infiniband, and also by 10 Gbps Ethernet. Most experiments use a single core of each server with the following configuration: 1 server as router, 7 servers in the processing tier, 4 servers in the storage tier; and communication over Infiniband with remote direct memory access (RDMA). Infiniband is a modern network technology, that is mainstream ([3, 19]) — it achieves low latency and high throughput, e.g., infiniband allows RDMA in a few microseconds.

To implement our storage tier, we use RAMCloud [30], which provides high throughput and very low read/write latency, in the order of 5-10 μs for every put/ get operation. It is able to achieve this efficiency because it keeps all stored values in memory as a distributed key-value store, where a key is hashed to determine on which server the corresponding key-value pair will be stored. RAMCloud also provides high memory utilization because of its log-structured design, and "continuous availability" by harnessing large scale to recover from system failure in a few seconds.

The graph is stored as an adjacency list — every node-id in the graph is the key, and the corresponding value is an array of its 1-hop neighbors (both in- and out-neighbors). The graph is partitioned across storage servers via RAMCloud's default and inexpensive hash partitioning scheme, MurmurHash3 over graph nodes.

While processing queries, it is common to access the adjacency lists of several nodes at the same time. For example, when computing 2-hop neighbors, the processor must request the neighborhood information for each node that is in the 1-hop neighborhood. There would be a large amount of overhead if each request is made independently over the network. Fortunately, RAMCloud offers a *multi read* method in its API that allows the client to make batch requests and reduce communication costs.

We use a limited amount of main memory (0∼4GB) as the cache of query processors. Our codes are implemented in C++.

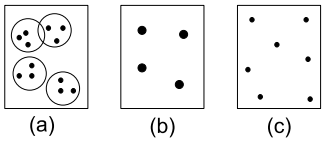

Figure 7: Three categories of query workloads: (a) r-hop hotspot, (b) concentrated hotspot, and (c) uniform workload.

• **Datasets.** We summarize our data sets in Table 1. As explained in Section 2, we store both in- and out-neighbors of a node.

WebGraph: This is the uk-2007-05 web graph downloaded from http://law.di.unimi.it/datasets.php. This graph is a collection of UK web pages, which are represented as nodes. If there is a hyperlink in page u to page v, we assign an edge between them.

Friendster: We downloaded this graph from https://snap.stanford.edu, which is a subgraph of the original *Friendster* social network. The nodes represent users and the edges denote friendship links.

Memetracker: This dataset tracks the quotes and phrases that appeared most frequently from August 1 to October 31, 2008 across the entire online news spectrum. We downloaded the dataset from https://snap.stanford.edu, and constructed a graph as follows. We consider the documents as nodes and hyper-links as edges.

Freebase: We downloaded the *Freebase* knowledge graph from http://www.freebase.com/. The nodes are named entities (e.g., Google) or abstract concepts (e.g., Asian people), and edges denote relations (e.g., founder) between two nodes.

• **Online Query Workloads.** We consider three classic online graph queries [53], all require traversals up to h hops: (1) *h-hop Neighbor Aggregation*: Count the number of h-hop neighbors of a query node. (2) *h-step Random Walk with Restart*: The query starts at a node, and at each successive step jumps to one of its neighbors with equal probability, or returns to the query node with a small probability. The query consists of h steps. (3) *h-hop Reachability*: Given a source and a target node, find if the target is reachable from the source within h-hops. In all our experiments, we consider a uniform mixture of the above queries. To simulate various real-world applications, we generate three categories of workloads by varying how the query nodes are selected, as illustrated in Figure 7.

r-hop Hotspot, h-hop Traversal: We simulate the scenario when queries are drawn from a hotspot region; and the hotspots change over time. In particular, we select 100 nodes from the graph uniformly at random. Then, for each of these nodes, we select 10 different query nodes which are at most r-hops away from that node. Thus, we generate 1000 queries; every 10 of them are from one hotspot region, and the pairwise distance between any two nodes from the same hotspot is at most $2r$. Finally, all queries from the same hotspot are grouped together and sent consecutively. We report our results averaged over these 1000 queries.

Concentrated Hotspot, h-hop Traversal: This is a specific case of the above workload, with $r = 0$. Thus, we select 100 query nodes from the graph uniformly at random. Each query node is sent 10 times consecutively, before the next set of queries. We report our results averaged over 1000 queries.

Uniform Workload, h-hop Traversal: We select 1000 query nodes uniformly at random, and report our results averaged over these queries. In contrast to the earlier hotspot-based workloads, it is less likely that the uniform workload will have query nodes which are neighbors of one another.

For the r-hop hostspot workload, in order to realize the effect of topology-aware locality, we consider smaller values of r: 1 and 2. We also vary h from 1∼4 in our experiments. Due to the power-law nature of real-world networks, all our graphs have small diameter.

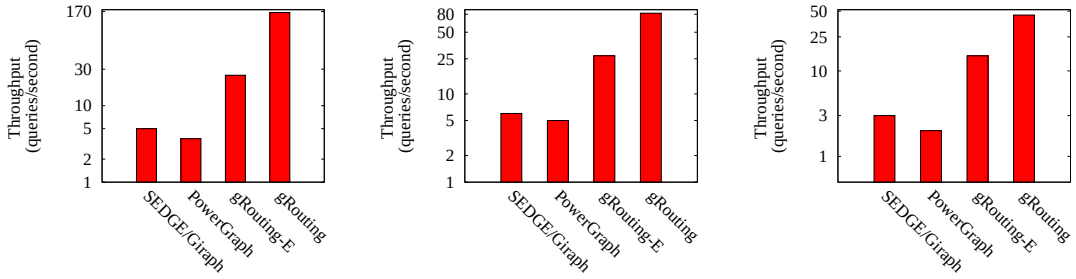

(a) *WebGraph* (b) *MemeTracker* (c) *Freebase*

Figure 8: Throughput comparison: 2-hop hotspot, 3-hop traversal workload. SEDGE/ Giraph and PowerGraph: 12 Servers. gRouting-E: our design with Ethernet. gRouting: our design with Infiniband. We use 7 query servers, 4 storage servers, and 1 router in our design.

Therefore, traversing beyond 4 hops accesses a significant portion of the whole graph, which is not very interesting from the perspective of ad-hoc online queries that we consider in this work.

• **Evaluation Metrics.** We measure performance as follows.

Query Response Time: This is the time required to answer one query. We report the response time averaged over all queries.

Query Processing Throughput: The throughput measures the number of queries that can be processed per unit time.

Cache Hit Rate: We report cache hit rates, since higher cache hit rates reduce the query response time. Consider t queries $q_1, q_2, \ldots, q_t$ received successively by the router. For simplicity, let us assume that each query retrieves all h-hop neighbors of that query node (i.e., h-hop neighborhood aggregation). We denote by $|N_h(q_i)|$ the number of nodes within h-hops from q_i. Among them, we assume that $|N_h^c(q_i)|$ number of nodes are found in the query processors' cache. We define cache hit and miss rates as follows.

$$\text{Cache Hit Rates} := \sum_{i=1}^{t} |N_h^c(q_i)|$$

$$\text{Cache Miss Rates} := \sum_{i=1}^{t} (|N_h(q_i)| - |N_h^c(q_i)|) \quad (8)$$

• **Parameter Setting.** We find that *embed routing performs the best compared to three other routing strategies*. We also set the following parameter values unless otherwise stated. We select these values since they perform the best in our implementation. We shall, however, demonstrate sensitivity of our routing algorithms with these parameters in the Appendix.

We use 4GB cache in each query processor. All experiments are performed with the cache initially empty (cold cache). The number of landmarks $|L|$ is set as 96 with at least 3 hops of separation from each other. For graph embedding, 10 dimensions are used. Load Factor is set as 20, and the smoothing parameter $\alpha = 0.5$.

In order to realize how our routing schemes perform as compared to the scenario when there is no cache in the query processors, we consider an additional "no-cache" scheme in our experiments. In this mode, all queries are routed following the next ready routing technique; however, as there is no cache in the query processors, there will be no overhead due to cache lookup and maintenance.

4.3 Comparison with State-of-the-Art

We compare gRouting with two distributed graph processing systems, SEDGE/Giraph [53] and PowerGraph [9]. These implementations are not compatible with Infiniband network and RDMA. Thus, for a fairer comparison, we consider a version of gRouting, completely on Ethernet (gRouting-E). We consider 12 machines configuration of SEDGE and PowerGraph, since query processing and graph storage in them are coupled on same machines. On the contrary, we fix the number of routing, processing, and storage servers as 1, 7 and 4, respectively, in our case.

In Figure 8, we show comparison results with 2-hop hotspot, 3-hop traversal queries. The average 3-hop neighborhood size varies from 100K$\sim$1000K nodes over these datasets. We find that our throughput, with hash partitioning, and over Ethernet, is 5$\sim$10 times better than two distributed graph processing systems (SEDGE and PowerGraph) with expensive graph partitioning and re-partitioning strategies. The graph re-partitioning in SEDGE requires around 1 hour and also apriori information on future queries, whereas PowerGraph graph partitioning finishes in about 30 min. On the contrary, gRouting performs lightweight hash partitioning over graph nodes, and does not require any prior knowledge of the future workloads. Moreover, our throughput over Infiniband is 10$\sim$35 higher than these systems. *These results demonstrate the usefulness of smart query routing over expensive graph partitioning and re-partitioning strategies*.

In the following, we present our scalability results, impact of cache sizes, and graph updates over *Webgraph* because this is the largest graph dataset that we consider in our experiments. We obtain similar results over other datasets; and hence, they are omitted for brevity. We emphasize on r-hop hotspot, h-hop traversal queries because this workload simulates a dynamic hotspot scenario. We shall, however, demonstrate the efficiency of our routing algorithms over other workloads and datasets in the Appendix.

4.4 Scalability and Deployment Flexibility

One of the main benefits of separating the processing and storage tiers in our design is deployment flexibility — they can be scaled-up independently. In these experiments, we investigate how our performance scales with respect to both these tiers.

Processing Tier: We vary the number of processing servers from 1 to 7, while using 1 router and 4 storage servers. The maximum of 7 processing servers are used, since the cluster on which we run our experiments has 12 machines. We consider 2-hop hotspot, 3-hop traversal queries over *Webgraph*. In Figure 9(a), we show throughput with varying number of processing servers. Corresponding cache hit rates are presented in Figure 9(b).

For these experiments, we assume that each query processor has sufficient cache capacity (4GB) to store the results of all 1000 queries (i.e, adjacency lists of 52M nodes, shown in Figure 9(b)). Since, for every experiment, we start with an empty cache, and then send the same 1000 queries in order, maximum cache hit happens when there is only one query processor. When we increase the number of query processors, these queries get distributed and processed by different processors, thus cache hit rate generally decreases. This is more evident for our baseline routing schemes, and

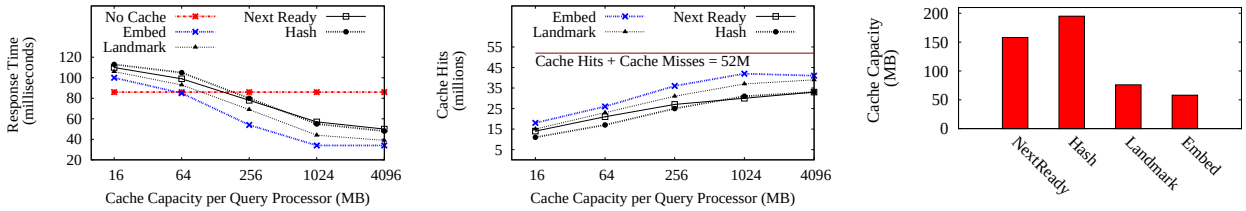

(a) Response time (b) Cache hits (c) Min. cache required to reach No-Cache's response time

Figure 11: Impact of cache size: *WebGraph*, 2-hop hotspot, 3-hop traversal workload

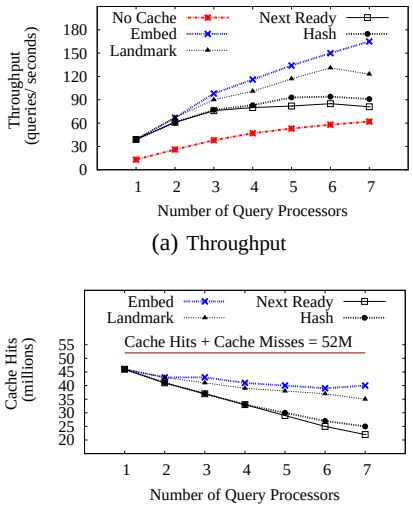

(a) Throughput

(b) Cache Hits

Figure 9: Performance with varying number of query processors: 2-hop hotspot, 3-hop traversal workload. *WebGraph*.

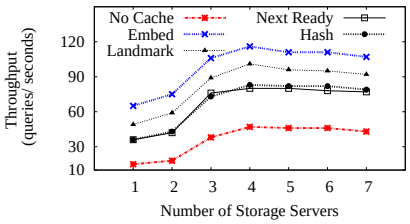

Figure 10: Performance with varying number of storage servers: 2-hop hotspot, 3-hop traversal workload. *WebGraph*.

we find that their throughput saturates at around 3∼5 query servers. These findings, in fact, demonstrate the usefulness of smart query routing, that is, in order to maintain a similar cache hit rate, the queries need to be routed intelligently in a distributed setting. As shown in Figure 9(b), *since Embed routing is able to sustain almost same cache hit rate with more number of query processors, its throughput scales linearly with the number of query processors*.

Storage Tier: In these experiments, we vary the number of storage servers from 1 to 7, whereas 1 server is used as the router and 4 servers as query processors (Figure 10). When we use 1 storage server, we are still able to load the entire 60GB *Webgraph* on the main memory of that server, since each of our cluster machines has sufficient RAM.

The throughput is the least when there is only one storage server. We observe that 1∼2 storage servers are insufficient to handle the demand created by 4 query processors. However, with 4 storage servers, the throughput saturates, since the bottleneck is transferred to query processors. This is evident from our previous results — the throughput with 4 query processors was about 120 queries per second (Figure 9(a)), which is the same throughput achieved with 4 storage servers in the current experiments. Therefore, if we keep the number of query processors fixed as 4, and employ more than 4 storage servers, throughput saturates and even slightly degrades with adding more storage. This happens due to the overhead of storage coordinator. We use one coordinator instance, which acts as the storage interface, dispatches requests to storage servers, and aggregates the retrieved results. Once there are more storage servers, this increases the overhead of the coordinator instance, causing it to spend more time per request.

Landmark routing	Embed routing	
	Embed landmarks	Embed per node
35 sec	36 sec	1 sec

Table 2: Preprocessing times, *WebGraph*

Landmark routing	Embed routing	Original graph size
2.8 GB	4GB	60.3GB

Table 3: Preprocessing storage, *WebGraph*

4.5 Impact of Cache Sizes

In previous experiments, we assign 4GB cache to each processor, which was large enough for our queries; and we never discarded anything from the cache. In reality, this might not always be feasible. Therefore, we perform experiments to verify how the system behaves when it needs to evict cache entries.

In Figure 11, we present average response times with various cache capacities. At the largest, with 4GB cache per processor, no eviction occurs. Therefore, there is no additional performance gain by increasing the cache capacity. On the other extreme, having cache with less than 64MB per processor results in worse response times than what was obtained with no-cache scheme, represented by the horizontal red line (86ms in Figure 11). When the cache does not have much space, it ends up evicting entries that might have been useful in the future. Hence, there are not enough cache hits to justify its maintenance and lookup costs when the cache size is smaller than 64MB per processor.

Using the same results, one can evaluate our routing strategies from the perspective of minimum cache requirement in order to achieve a response time of 86ms, the break-even point of deciding whether or not to add a cache. We can observe in Figure 11(c) that smart routing schemes are able to achieve this response time with a much lower cache requirement, as compared to that of the baselines. These results illustrate that *our smart routings utilize the cache space well; and for the same amount of cache space, they achieve lower response time compared to baseline routings*.

4.6 Preprocessing and Graph Updates

Preprocessing Time and Storage: We present gRouting preprocessing times in Table 2. For landmarks routing, one needs to compute the distance of every node to all landmarks, which can be evaluated by performing a BFS from each landmark over the whole

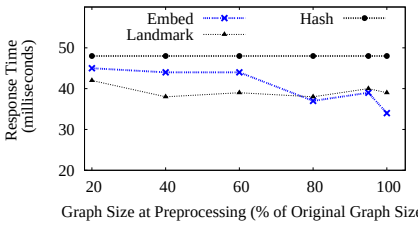

Figure 12: Robustness of landmark and embed routings with respect to dynamic graph updates: *WebGraph*, 2-hop hotspot, 3-hop traversal

graph. This takes about 35 sec for one landmark in *Webgraph*, and can be parallelized per landmark. For embed routing, in addition, we need to embed every node with respect to landmarks, which requires about 1 sec per node in *Webgraph*, and is again parallelizable per node.

The preprocessed information requires $\mathcal{O}(nP)$ space for landmark routing, where n is the number of nodes and P the number of processors. In our implementation, each integer value consumes 4-bytes. Therefore, the preprocessed landmark routing information consumes about 2.8GB storage space in case of *Webgraph* and with 7 query processors. On the contrary, our embed routing requires $\mathcal{O}(nD)$ space for storing the preprocessed information, where n is the number of nodes and D the number of dimensions. For *Webgraph*, with embedding dimensionality 10, the graph embedding size is only 4GB. Both these preprocessed information are modest compared to the original *Webgraph* size, which is around 60GB (Table 3).

Graph Updates: In these experiments, we perform minimal changes (often none) to our preprocessed information with every graph update, and analyze how robust our routing schemes are with respect to graph modification.

In particular, we use *Webgraph* due to its larger size, and then perform our preprocessing step at a reduced subgraph of this dataset. For example, at 20% of the original dataset (Figure 12), we select only 20% of all nodes uniformly at random, and compute preprocessed information over the subgraph induced by these selected nodes. However, we always run our query over the complete *Webgraph*. We incrementally compute the necessary information for the new nodes, as they are being added, without changing anything on the preprocessed information of the earlier nodes. For an example, in case of embed routing, we only compute the distance of a new node to the landmarks, and thereby find the coordinates of that new node. However, one may note that with the addition of every new node and its adjacent edges, the preprocessed information becomes outdated (e.g., the distance between two earlier nodes might decrease). Since we do not change anything on the preprocessed information, this experiment demonstrates the robustness of our method with respect to graph updates.

Figure 12 depicts that *our smart routing schemes are robust for a small number of graph updates*. With embed routing, preprocessed information over the whole graph results in response time of 34 ms, whereas preprocessed information at 80% of the graph results in response time of 37 ms (i.e., response time increases by only 3 ms). As expected, the response time deteriorates when preprocessing is performed on a smaller amount of graph data, e.g., with only 20% graph data, response time increases to 44 ms, which is comparable to the response time of baseline hash routing (48 ms).

Landmark routing is more robust — we do not find any significant difference in the response time even when the preprocessing is done at 40% of the original graph. This is because, for a reduced-size graph, we select the nodes *uniformly at random* from the original graph. Hence, the landmarks obtained from the reduced-sized graph also retains their effectiveness over the original graph. This explains the robustness of landmark routing. Nevertheless, for significantly reduced-sized graphs (e.g., preprocessing at only 20% of the original graph), the response time of landmark routing also increases to 42 ms.

We note that the incremental update time is quite small. Addition of a new node requires computing its distance to all landmarks, which has the same time complexity of performing a breadth first search (BFS) from the node (35 sec for *Webgraph* dataset). Then, we embed this node with respect to landmarks (1 sec for each node in *Webgraph*, and can be parallelized for every node). In essence, one can perform such incremental updates in a batch manner.

4.7 Summary of Experimental Results

We summarize our experimental results as follows. **(1)** gRouting, with lightweight hash partitioning, achieves 5∼10 times more throughput than two state-of-the-art distributed graph processing systems: SEDGE/Giraph and PowerGraph, with expensive graph partitioning and re-partitioning. Moreover, gRouting's throughput over Infiniband is 10∼35 higher than these systems. **(2)** Our smart routing schemes, landmark and embed, achieves more throughput and less query response time compared to baseline routing schemes, e.g., no-cache, next ready, and hash. **(3)** gRouting's throughput increases linearly with more number of query processors. **(4)** gRouting is robust with respect to workload changes and graph updates.

5. RELATED WORK

We studied smart query routing for distributed graph querying — a problem for which we are not aware of any prior solution. In the following we, however, provide a brief overview of work in neighborhood areas.

Landmarks and Graph Embedding. Kleinberg et. al. [17] discussed the problem of approximating graph distances using a small set of beacons (i.e., landmarks). Landmarks were also used in path finding, shortest path estimation, as well as in estimating network properties [1, 34, 35]. Graph embedding algorithms [56], on the other hand, place nodes at points on some surface such that the inherent graph structure is preserved. Traditionally, spectral methods such as principal component analysis (PCA) have been applied to many graph embedding and dimensionality reduction tasks, where the underling manifold is a linear subspace. For non-linear cases, many recent techniques, e.g., multidimensional scaling (MDS), locally linear embedding (LLE), Laplacian eigenmaps (LEM), Isomap, semidefinite embedding (SDE), minimum volume embedding (MVE), and structure preserving embedding (SPE) are proposed. In the past, graph embedding schemes were effectively employed in internet routing, such as predicting internet network distances and estimating minimum round trip time between hosts [6]. To the best of our knowledge, ours is the first study that applies graph embedding and landmark techniques to design effective routing algorithms for distributed graph querying.

Graph Partitioning, Re-partitioning, and Replication. The balanced, minimum-edge-cut graph partitioning divides a graph into k partitions such that each partition contains about the same number of nodes (and/or edges), and the number of cut-edges is minimized. Even for $k = 2$, the problem is **NP**-hard, and there is no approximation algorithm with a constant approximation ratio unless **P** =**NP** [28]. Therefore, efforts were made in developing polynomial-time heuristics — METIS, Chaco, SCOTCH, to name a few. Their parallel versions were also proposed, e.g., ParMETIS. However, real-world graphs, due to power-law degree

distribution, are difficult to partition in a balanced minimum edge-cut manner. PowerGraph [9] mitigates this problem by performing a balanced node-cut of the graph. More sophisticated graph partitioning schemes were also proposed, e.g., complementary partitioning [53], multi-level [49] and balanced [47] label propagation, and network-bandwidth-aware graph partitioning [5].

Graph re-partitioning, on the other hand, is critical for (1) online queries, since workloads on nodes change frequently [23, 40], and for (2) dynamic graphs where nodes and edges are added/ deleted. Under the first category, SEDGE [53] and Mizan [16] developed graph re-partitioning techniques based on past workloads. Nicoara et. al. [28] proposed light-weight graph re-partitioning strategies by taking into account the data access patterns of previous queries. Under the second category, incremental (e.g., one-pass) and re-partitioning techniques were proposed for dynamic [48, 51, 57] and stream graphs [44]. However, online monitoring of the load imbalance and graph updates, dynamically modifying the graph partitions, and migrating data across servers are expensive. As discussed earlier, with our embed routing, we are able to bypass these expensive graph partitioning and re-partitioning problems to the existing cache replacement policy of the query processors.

Replication has been used for graph partitioning, re-partitioning, dynamic load balancing, and fault tolerance. In earlier works, [11, 33] proposed one extreme version of it by replicating the graph sufficiently so that, for every node in the graph, all of its neighbors are present locally. Mondal et. al. designed an overlapping graph re-partitioning scheme [23], which updates its partitions based on the past read/write patterns. Very recently, Huang et. al. [10] designed a lightweight re-partitioning and replication scheme considering access locality, fault tolerance, and dynamic graph updates. We note that we also replicate the graph data at query processors' cache and there could be overlaps between two processors' cache contents. However, we only replicate the active regions of the graph based on the most recent workloads, and more importantly, unlike [10, 23] we do not explicitly run any graph re-partitioning or replication strategy at our processors or storage servers. Instead, our smart routing algorithms automatically perform such effective and dynamic replications at the processors' cache.

Graph Caching, Views, De-coupling, and Multi-Query Optimization. Facebook uses a fast caching layer, Memcached on top of a graph database to scale the performance of graph querying [29]. Recently, graph-structure-aware and workload-adaptive caching techniques were proposed, e.g., [2, 31]. We use the LRU (i.e., least recently used) replacement policy in the query processors' cache for simplicity, and it works well together with our smart query routing. There are other works on view-based graph query answering [8] and multi-query optimizations [18]. Unlike ours, these approaches require the workload to be known in advance.

Burkhardt et. al. implemented BFS on top of Map-Reduce, together with a de-coupled key-value based graph storage for better adjacency [4]. Our proposed smart routing logics are orthogonal to them. Very recently, Shalita et. al. [39] employed decoupling for an optimal assignment of HTTP requests over a distributed graph storage. First, they perform a static partitioning of the graph data in storage servers based on graph structure and co-access patterns. Next, they find the past workload on each partition, and dynamically assign these partitions to query processors such that load balancing can be achieved. While their decoupling principle and dynamic assignment at query processors are similar to ours, they still explicitly perform a sophisticated graph partitioning at the storage servers, and update such partitions in an offline manner. On the contrary, our smart routing algorithms, as reasoned earlier, automatically partition the active regions of the graph in a dynamic manner and store them in the query processors' cache, thereby achieving both load balancing and improved cache hit rates.

Large-Scale Graph Processing Systems. Based on parallel architecture, large-scale graph processing systems can be categorized into distributed-memory (e.g., Pregel [21], PowerGraph [9]) and shared-memory (Ligra [42], Galios [27]) systems. While shared-memory systems usually have lower communication cost, distributed-memory systems scale well.

The existing graph processing systems can also be classified based on their application support: **(1)** systems for offline graph analytics and **(2)** systems for online graph querying. Offline graph analytic systems perform iterative, batch processing over the graph by following a node/edge-centric scatter gather model. Examples of such computations include PageRank, diameter of a graph, and strongly connected components. Pregel [21] and PowerGraph [9] are representatives of offline graph processing systems. For a survey, we refer to [52].

Offline graph processing systems are often not suitable for online graph queries [41, 55]. This is because online graph querying requires fast response time, and these queries explore parts of the graph. Examples of online graph queries are h-hop neighborhood search, shortest-path finding, sub-graph matching, and SPARQL queries. As discussed in [41], online graph querying systems must facilitate graph traversal over a region of the graph for faster query answering. Horton [37] and Ligra [42] are examples of online graph querying systems. Trinity [41] supports both offline graph analytics and online graph queries. The de-coupled architecture and our designed smart routing algorithms, being generic, can benefit the existing online graph querying systems.

NoSQL and Graph Data Stores. NoSQL databases are usually structured as trees, graphs, XML, RDF, or key-value stores. They aim at simplicity, horizontal scaling, availability, and partition tolerance — often at the cost of consistency. An incomplete list of today's NoSQL databases include **(1)** key-value stores: Redis, RAMCloud; **(2)** document stores: MongoDB, Couchbase; (3) graph databases: Neo4J (currently does not support distributed graph storage), InfiniteGraph. Several distributed SPARQL engines were also designed, e.g., RDF-3X [26] for processing RDF data. Many of these data stores can be used as our distributed storage layer. We use RAMCloud [30] in our implementation due to its low latency, high availability, and high memory utilization.

6. CONCLUSIONS

We studied h-hop traversal queries – a generalized form of various online graph queries that access a small region of the graph, and require fast response time. To answer such queries with low latency and high throughput, we follow the principle of decoupling query processors from graph storage. Our work emphasized *less* on the requirements for an expensive graph partitioning and re-partitioning technique, instead we developed smart query routing strategies for effectively leveraging the query processors' cache contents, thereby improving the throughput and reducing latency of distributed graph querying. In addition to workload balancing, deployment flexibility, and fault tolerance, gRouting is able to provide linear scalability in throughput with more number of query processors, works well in the presence of query hotspots, and is also adaptive to workload changes and graph updates.

Future work can be in two directions. First, enabling modern query processing techniques and hardware trends, e.g., parallel, collaborative scans in main memory could be interesting to improve query throughput and to support a larger set of queries. Second, it would be worth investigating transaction workloads such as stream-

ing graph updates and consistency criteria in a decoupled system.

7. REFERENCES

[1] T. Akiba, Y. Iwata, and Y. Yoshida. Fast Exact Shortest-path Distance Queries on Large Networks by Pruned Landmark Labeling. In *SIGMOD*, 2013.

[2] H. Aksu, M. Canim, Y. Chang, I. Korpeoglu, and Ö. Ulusoy. Graph Aware Caching Policy for Distributed Graph Stores. In *IC2E*, 2015.

[3] C. Binnig, A. Crotty, A. Galakatos, T. Kraska, and E. Zamanian. The End of Slow Networks: It's Time for a Redesign [Vision]. In *VLDB*, 2016.

[4] P. Burkhardt and C. A. Waring. A Cloud-Based Approach to Big Graphs. In *HPEC*, 2015.

[5] R. Chen, M. Yang, X. Weng, B. Choi, B. He, and X. Li. Improving Large Graph Processing on Partitioned Graphs in the Cloud. In *SoCC*, 2012.

[6] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A Decentralized Network Coordinate System. In *SIGCOMM*, 2004.

[7] J. Dinan, D. B. Larkins, P. Sadayappan, S. Krishnamoorthy, and J. Nieplocha. Scalable Work Stealing. In *SC*, 2009.

[8] W. Fan, X. Wang, and Y. Wu. Answering Graph Pattern Queries using Views. In *ICDE*, 2014.

[9] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin. PowerGraph: Distributed Graph-parallel Computation on Natural Graphs. In *OSDI*, 2012.

[10] J. Huang and D. J. Abadi. Leopard: Lightweight Edge-oriented Partitioning and Replication for Dynamic Graphs. In *VLDB*, 2016.

[11] J. Huang, D. J. Abadi, and K. Ren. Scalable SPARQL Querying of Large RDF Graphs. In *VLDB*, 2011.

[12] G. Jeh and J. Widom. Scaling Personalized Web Search. In *WWW*, 2003.

[13] R. Jin, H. Hong, H. Wang, N. Ruan, and Y. Xiang. Computing Label-Constraint Reachability in Graph Databases. In *SIGMOD*, 2010.

[14] R. Jin, L. Liu, B. Ding, and H. Wang. Distance-Constraint Reachability Computation in Uncertain Graphs. In *VLDB*, 2011.

[15] G. Karypis. METIS and ParMETIS. In *Encyclopedia of Parallel Computing*. Springer, 2011.

[16] Z. Khayyat, K. Awara, A. Alonazi, H. Jamjoom, D. Williams, and P. Kalnis. Mizan: A System for Dynamic Load Balancing in Large-scale Graph Processing. In *EuroSys*, 2013.

[17] J. Kleinberg, A. Slivkins, and T. Wexler. Triangulation and Embedding Using Small Sets of Beacons. *J. ACM*, 56(6):32:1–32:37, 2009.

[18] W. Le, A. Kementsietsidis, S. Duan, and F. Li. Scalable Multi-query Optimization for SPARQL. In *ICDE*, 2012.

[19] S. Loesing, M. Pilman, T. Etter, and D. Kossmann. On the Design and Scalability of Distributed Shared-Data Databases. In *SIGMOD*, 2015.

[20] K. Macropol, T. Can, and A. K. Singh. RRW: Repeated Random Walks on Genome-Scale Protein Networks for Local Cluster Discovery. *BMC Bioinformatics*, 10:283, 2009.

[21] G. Malewicz, M. H. Austern, A. J. C. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: A System for Large-scale Graph Processing. In *SIGMOD*, 2010.

[22] C. Martella, D. Logothetis, and G. Siganos. Spinner: Scalable Graph Partitioning for the Cloud. *CoRR*, abs/1404.3861, 2014.

[23] J. Mondal and A. Deshpande. Managing Large Dynamic Graphs Efficiently. In *SIGMOD*, 2012.

[24] J. Mondal and A. Deshpande. EAGr: Supporting Continuous Ego-centric Aggregate Queries over Large Dynamic Graphs. In *SIGMOD*, 2014.

[25] B. Nam, M. Shin, H. Andrade, and A. Sussman. Multiple Query Scheduling for Distributed Semantic Caches. *J. Parallel Distrib. Comput.*, 70(5):598–611, 2010.

[26] T. Neumann and G. Weikum. RDF-3X: a RISC-Style Engine for RDF. In *VLDB*, 2008.

[27] D. Nguyen, A. Lenharth, and K. Pingali. A Lightweight Infrastructure for Graph Analytics. In *SOSP*, 2013.

[28] D. Nicoara, S. Kamali, K. Daudjee, and L. Chen. Hermes: Dynamic Partitioning for Distributed Social Network Graph Databases. In *EDBT*, 2015.

[29] R. Nishtala, H. Fugal, S. Grimm, M. Kwiatkowski, H. Lee, H. C. Li, R. McElroy, M. Paleczny, D. Peek, P. Saab, D. Stafford, T. Tung, and V. Venkataramani. Scaling Memcache at Facebook. In *NSDI*, 2013.

[30] J. Ousterhout, P. Agrawal, D. Erickson, C. Kozyrakis, J. Leverich, D. Mazières, S. Mitra, A. Narayanan, G. Parulkar, M. Rosenblum, S. M. Rumble, E. Stratmann, and R. Stutsman. The Case for RAMClouds: Scalable High-performance Storage Entirely in DRAM. *SIGOPS Oper. Syst. Rev.*, 43(4):92–105, 2010.

[31] N. Papailiou, D. Tsoumakos, P. Karras, and N. Koziris. Graph-Aware, Workload-Adaptive SPARQL Query Caching. In *SIGMOD*, 2015.

[32] M. Potamias, F. Bonchi, C. Castillo, and A. Gionis. Fast Shortest Path Distance Estimation in Large Networks. In *CIKM*, 2009.

[33] J. M. Pujol, V. Erramilli, G. Siganos, X. Yang, N. Laoutaris, P. Chhabra, and P. Rodriguez. The Little Engine(s) That Could: Scaling Online Social Networks. In *SIGCOMM*, 2010.

[34] M. Qiao, H. Cheng, L. Chang, and J. X. Yu. Approximate Shortest Distance Computing: A Query-Dependent Local Landmark Scheme. In *ICDE*, 2012.

[35] M. J. Rattigan, M. E. Maier, and D. Jensen. Using Structure Indices for Efficient Approximation of Network Properties. In *KDD*, 2006.

[36] A. Roy, L. Bindschaelder, J. Malicevic, and W. Zwaenepoel. Chaos: Scale-out Graph Processing from Secondary Storage. In *SOSP*, 2015.

[37] M. Sarwat, S. Elnikety, Y. He, and M. F. Mokbel. Horton+: A Distributed System for Processing Declarative Reachability Queries over Partitioned Graphs. In *VLDB*, 2013.

[38] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad. Collective Classification in Network Data. *AI Magazine*, 29(3):93–106, 2008.

[39] A. Shalita, B. Karrer, I. Kabiljo, A. Sharma, A. Presta, A. Adcock, H. Kllapi, and M. Stumm. Social Hash: An Assignment Framework for Optimizing Distributed Systems Operations on Social Networks. In *NSDI*, 2016.

[40] Z. Shang and J. X. Yu. Catch the Wind: Graph Workload Balancing on Cloud. In *ICDE*, 2013.

[41] B. Shao, H. Wang, and Y. Li. Trinity: A Distributed Graph Engine on a Memory Cloud. In *SIGMOD*, 2013.

[42] J. Shun and G. E. Blelloch. Ligra: A Lightweight Graph Processing Framework for Shared Memory. In *PPoPP*, 2013.

[43] J. Shute, R. Vingralek, B. Samwel, B. Handy, C. Whipkey, E. Rollins, M. Oancea, K. Littlefield, D. Menestrina, S. Ellner, J. Cieslewicz, I. Rae, T. Stancescu, and H. Apte. F1: A Distributed SQL Database That Scales. In *VLDB*, 2013.

[44] I. Stanton and G. Kliot. Streaming Graph Partitioning for Large Distributed Graphs. In *KDD*, 2012.

[45] H. Tong, C. Faloutsos, B. Gallagher, and T. E.-Rad. Fast Best-effort Pattern Matching in Large Attributed Graphs. In *KDD*, 2007.

[46] K. Tretyakov, A. A.-Cervantes, L. G.-B. nuelos, J. Vilo, and M. Dumas. Fast Fully Dynamic Landmark-based Estimation of Shortest Path Distances in Very Large Graphs. In *CIKM*, 2011.

[47] J. Ugander and L. Backstrom. Balanced Label Propagation for Partitioning Massive Graphs. In *WSDM*, 2013.

[48] L. Vaquero, F. Cuadrado, D. Logothetis, and C. Martella. Adaptive Partitioning for Large-scale Dynamic Graphs. In *SOCC*, 2013.

[49] L. Wang, Y. Xiao, B. Shao, and H. Wang. How to Partition a Billion-Node Graph. In *ICDE*, 2014.

[50] J. Wei, F. Xia, C. Sha, C. Xu, X. He, and A. Zhou. Workload-Aware Cache for Social Media Data. In *APWeb*, 2013.

[51] N. Xu, L. Chen, and B. Cui. LogGP: A Log-based Dynamic Graph Partitioning Method. In *VLDB*, 2014.

[52] D. Yan, Y. Bu, Y. Tian, A. Deshpande, and J. Cheng. Big Graph Analytics Systems. In *SIGMOD*, 2016.

[53] S. Yang, X. Yan, B. Zong, and A. Khan. Towards Effective Partition Management for Large Graphs. In *SIGMOD*, 2012.

[54] H. Zhang, G. Chen, B. C. Ooi, K. L. Tan, and M. Zhang. In-Memory Big Data Management and Processing: A Survey. *TKDE*, 27(7):1920–1948, 2015.

[55] Q. Zhang, D. Yan, and J. Cheng. Quegel: A General-Purpose System for Querying Big Graphs. In *SIGMOD*, 2016.

[56] X. Zhao, A. Sala, C. Wilson, H. Zheng, and B. Y. Zhao. Orion: Shortest Path Estimation for Large Social Graphs. In *WOSN*, 2010.

[57] A. Zheng, A. Labrinidis, and P. K. Chrysanthis. Planar: Parallel Lightweight Architecture-Aware Adaptive Graph Re-partitioning. In *ICDE*, pages 121–132, 2016.

APPENDIX

A. SENSITIVITY ANALYSIS

We analyze the impact of various parameters including the number of landmarks, their minimum separation, embedding dimensionality, load factor, and smoothing factor (α) on gRouting's performance. We find that gRouting is more sensitive towards load factor and embedding dimensionality. We demonstrate our results with 2-hop hotspot, 3-hop traversal queries over *Webgraph*. In all Figures, we also present the response time (or, throughput) of our best baseline, which is hash routing, for an effective comparison.

Load Factor: This parameter impacts both of our smart routing schemes. We find from Equations 3 and 7 that smaller values of load factor diminish the impact of "smart" routing (i.e., landmarks and node co-ordinates), instead queries will be routed to the processor having the minimum workload. On the other hand, higher values of load factor reduces the impact of load balancing (i.e., query stealing) — queries would be routed solely based on landmarks and node co-ordinates. Therefore, in these experiments, we expect that

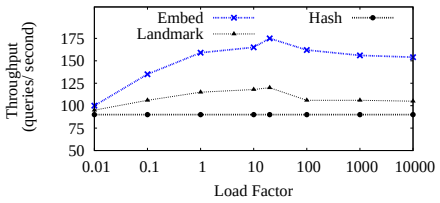

Figure 13: Impact of load factor: *WebGraph*, 2-hop hotspot, 3-hop traversal workload

the throughput will initially increase with higher values of load factor, until it reaches a maximum, and then it would start decreasing. Indeed, it can be observed in Figure 13 that with load factor between 10∼20, the best throughput is achieved for both smart routing schemes. As stated earlier, when the load factor is very large, e.g., 10000, the load balancing influence in the smart routing tends to be zero. Hence, the measurements at 10000 indicates what our throughput would have been if load balancing (via query stealing) was not implemented.

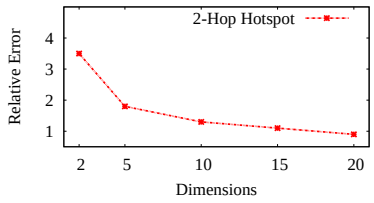

(a) Distance error between node pairs

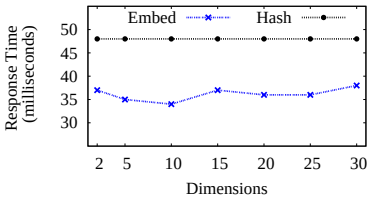

(b) Response time

Figure 14: Impact of embedding dimensionality: *WebGraph*, 2-hop hotspot, 3-hop traversal workload

Embedding Dimensionality: We consider the performance implications of the number of dimensions for our embed routing. For these experiments, we create several embeddings, with dimensionality from 2 to 30. While the relative error in distance between node pairs decreases with higher dimensions, it almost saturates after 10 dimensions (Figure 14(a)). On the other hand, we observe that the average response time reduces until dimension 10, and then it slowly increases with more dimensions (Figure 14(b)). This is because with higher dimensions, we reduce the distance prediction error, thereby correctly routing the queries and getting more cache hits. However, a large number of dimensions also increases the routing decision making time at the router. Hence, the least response time is achieved at dimensionality 10.

Number of Landmarks: Figure 15(a) shows the variation of query response time with the number of landmark nodes. Both our smart routing strategies, i.e., landmark and embed get affected by the number of landmarks. For landmark routing, the first P pivot landmarks are selected as those landmarks which are farthest apart from each other (see Algorithm 1), and then each pivot landmark is assigned to one distinct query processor. Since P, i.e., the number of query processors is quite small, usually a small number of initial landmarks (e.g., 32 in Figure 15(a)) would be sufficient to find

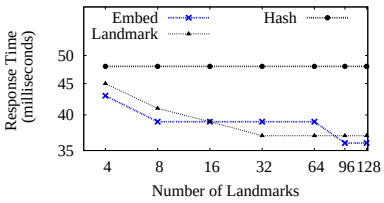

(a) Response time

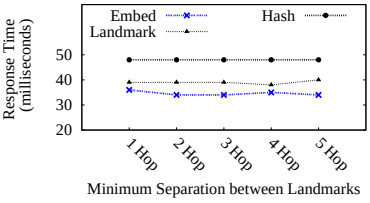

(b) Response time

Figure 15: Impact of landmark numbers and their minimum separation: *WebGraph*, 2-hop hotspot, 3-hop traversal workload

these P high-quality pivot landmarks.

On the other hand, the quality of graph embedding depends on the number of landmarks selected, generally the more, the better (Algorithm 3). we find that at 96 landmarks, there is a significant reduction in response time for embed routing. We also recall that, in the preprocessing phase of our smart routings, one needs to compute the distance of every landmark to all nodes, and thus, preprocessing time increases with more landmarks. Hence, we perform a trade-off between query response time and preprocessing time, thereby setting the optimal number of landmarks as 96 in our experiments.

Minimum Separation between Landmark Pairs: Next, we verify how the minimum separation between landmark pairs ("threshold" on line 4, Algorithm 1) affects the average response time. We find that the separation does not have a significant influence on response time, with embed's best performance at 3-hops separation and landmark's at 4-hops separation (Figure 15(b)).

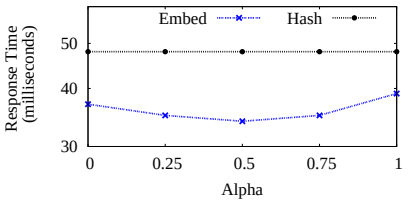

Figure 16: Impact of smoothing parameter (α): *WebGraph*, 2-hop hotspot, 3-hop traversal workload

Smoothing Parameter: The smoothing parameter (α) effects only embed routing, and it can take values from 0 to 1, indicating how much new queries affect the query processor's expected moving average (EMA). A value of 1 would mean that the processor's EMA would be equal to the last query's co-ordinates, while a value of 0 would cause the processor's EMA to retain its initial value. We find in Figure 16 that the response time for embed routing scheme reduces when α is in the range of (0.25, 0.75).

B. EXPERIMENTS WITH OTHER WORKLOADS AND DATASETS

B.1 Efficiency with Different Workloads

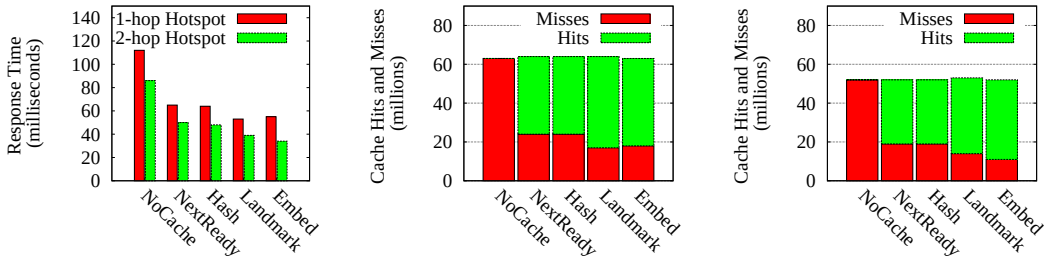

(a) Response time (b) Cache hits and misses (1-hop hotspot) (c) Cache hits and misses (2-hop hotspot)

Figure 17: Efficiency: r-hop hotspot, 3-hop traversal workload. We vary r = 1, 2. *WebGraph*

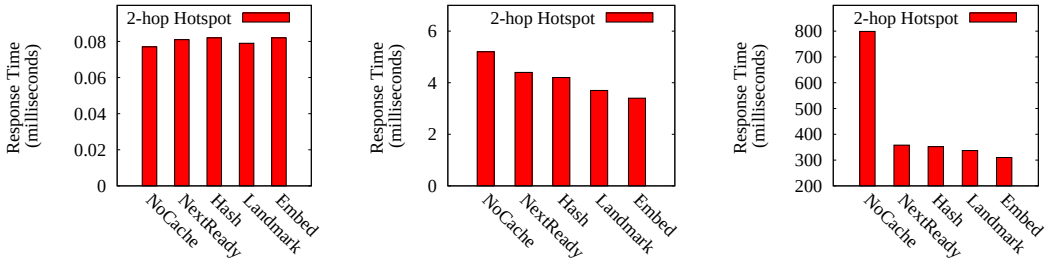

(a) 1-hop traversal (b) 2-hop traversal (c) 4-hop traversal

Figure 18: Efficiency: 2-hop hotspot, h-hop traversal workload. We vary h = 1∼4. *WebGraph*.

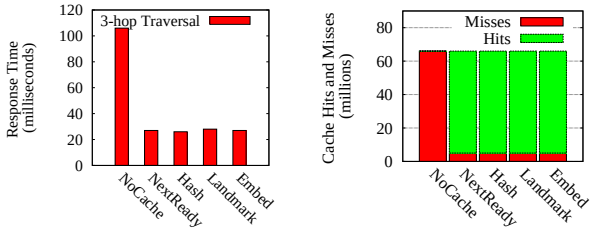

(a) Response time (b) Cache hits and misses

Figure 19: Efficiency: concentrated hotspot, 3-hop traversal. *WebGraph*.

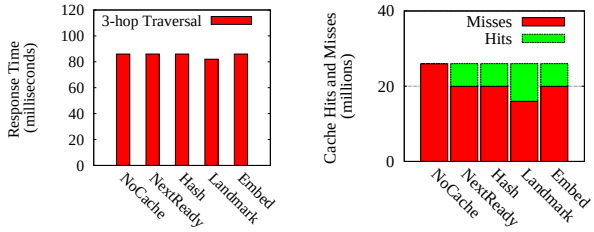

(a) Response time (b) Cache hits and misses

Figure 20: Efficiency: uniform, 3-hop traversal. *WebGraph*.

We first analyze the efficiency for r-hop hotspot, h-hop traversal workloads. In Figure 17, we vary r as 1 and 2, while setting h as 3. The average response times over *Webgraph* are shown in Figure 17(a), and the corresponding numbers of cache hits and misses are illustrated in Figures 17(b) and 17(c).

We find that adding cache to the processors for this category of workloads reduces the average response time significantly. Even baseline routing schemes such as next ready and hash obtain up to a 44% reduction in the average response time (86ms for no-cache vs. 48ms for hash routing). In addition, we observe that our smart routing techniques, landmark and embed, both outperform the baseline routing algorithms by a margin of 29% (48ms for hash vs. 34ms for embed). This is because our smart routing schemes are designed to capture topology-aware locality; and hence, make better use of the processors' cache contents. In Figures 17(b) and 17(c), one may note that the smart routing schemes are indeed obtaining more cache hits as compared to that of baseline routings. Finally, we analyze the difference in average response times for 1-hop and 2-hop hotspots in Figure 17(a). It happens because these are two different workloads — the average 3-hop neighborhood size per query for our 1-hop hotspot workload is 717K nodes, while for the latter, the average 3-hop neighborhood size is 367K nodes.

Next, we run our experiments by setting r=2 and varying h as 1, 2 and 4. As depicted in Figure 18, the results for 2-hop traversal and 4-hop traversal are exactly similar to what we found earlier for 3-hop traversal; however, 1-hop traversal workload performs quite differently. For this workload (Figure 18(a)), all our routing schemes perform similarly, and the no-cache configuration is actually performing just as well or slightly better than other routing techniques. This is because for 1-hop traversal queries, the cache or the decoupled storage needs to be checked only once per query, since the result of a query node is its immediate 1-hop neighbors. Hence, there would be a cache hit only if the same query is requested twice, which is very unlikely to happen for our 2-hop hotspot workload. As a consequence, the cache does not contain any useful data to be leveraged. Moreover, all our routing schemes have some overhead due to cache maintenance and lookup. This explains why no-cache's average response time is slightly better for the 1-hop traversal workload.

For 4-hop traversals (Figure 18(c)), the difference between smart routing schemes and baselines has diminished compared to that for 2-hop and 3-hop traversal queries. Smart routing schemes still obtain more cache hits per query and have lower response times, e.g., 15% lower than baselines. However, due to 4-hop traversals, these queries process a larger amount of data. For example, the average 3-hop neighborhood size is 367K nodes, whereas the average 4-hop

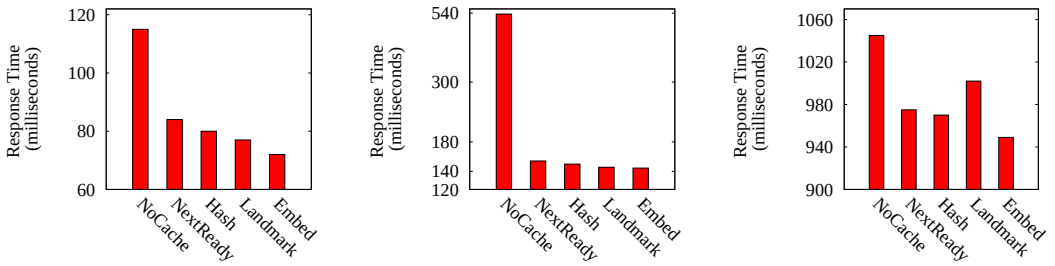

(a) *Memetracker* (b) *Freebase* (c) *Friendster*

Figure 21: Efficiency: *Memetracker*, *Freebase*, *Friendster*, 2-hop hotspot, 3-hop traversal workload

neighborhood size is 2M nodes. This means that computation occupies a significant part of the overall response time, reducing the relative impact of smart routing obtained via cache hits.

We demonstrate the performance of concentrated hotspot workloads in Figure 19. Due to repeated querying of same nodes, caching has even greater advantage. We find that all our routing schemes have a large reduction of up to 75% in response time when compared to the no-cache technique. However, the baseline ones perform almost as good as our smart routing schemes. This is because there is not much topology-aware locality that our smart routing schemes could benefit from, and the baselines perform equally well due to their load balancing techniques.

Finally, we evaluate our performance when the queries are chosen uniformly at random (Figure 20). As expected, cache is not very useful here — no-caching scheme has a similar response time as that of all our routing schemes except the landmark routing. Since landmark routing routes the queries based on their distances to landmarks, it is still able to route nearby queries to the same processor, even if those queries do not come consecutively. In particular, landmark routing improves the response time by 5% compared to other methods for this workload.

B.2 Efficiency over Different Datasets

We investigate how gRouting performs over other graphs, e.g., *Memetracker*, *Freebase*, and *Friendster*. We present the average response time with 2-hop hotspot, 3-hop traversal queries in Figure 21. The results with *Memetracker* are similar to those of our earlier results over *Webgraph*. In particular, baseline routing techniques reduce the response time by 30% compared to no-cache scheme, and smart routing algorithms further reduce it by 10% in comparison with baseline routings (Figure 21(a)).

Response times over *Freebase* are shown in Figure 21(b). One may note that all our routing schemes improve the response time by 70% as compared to no-cache strategy; however, there is no significant difference between smart routing schemes and baselines. In order to understand this behavior, we analyze the structure of *Freebase*, and find that due to smaller size of this graph, queries from different hotspots have significant overlap in their 3-hop neighborhoods, making all caching schemes equally beneficial towards reducing the average response time.

In case of *Friendster* (Figure 21(c)), baseline routings reduce the response time by 7% compared to no-cache scheme, and smart routing techniques reduce it by another 3% in comparison with baseline routings. Our investigation further reveals that the average 3-hop neighborhood size is 2.9M for this dataset, compared to 367K in case of *Webgraph*. Therefore, computation at the processing tier occupies a significant fraction of the overall response time. Besides, we find that the overlap across 3-hop neighborhoods for queries from the same hotspot region is lower than that of *Webgraph*. Thus caching becomes less effective here. Nevertheless, our embed smart routing reduces the response time by 40 ms (about 3%) as compared to baseline routing schemes.

Algorithm 5 h-hop Query Answering in gRouting: BFS Combined with Caching and Batch Requests to Storage

Require: Directed graph G, query node q, hop count h
Ensure: All nodes within h-hop of query node u

```
 1: nextSet ← φ, requestSet ← φ
 2: resultSet ← {u}, currentSet ← {u}
 3: hop = 1
 4: while hop ≤ h do
 5:    for all cur ∈ currentSet do
 6:       if cur in cache then
 7:          for all next ∈ Out-Neighbor(cur) do
 8:             if next ∉ resultSet then
 9:                resultSet ← resultSet ∪ {next}
10:                nextSet ← nextSet ∪ {next}
11:             end if
12:          end for
13:       else
14:          requestSet ← requestSet ∪ {cur}
15:       end if
16:    end for
17:    if requestSet ≠ φ then
18:       execute batch request to storage: multi-read(requestSet)
19:       for all cur ∈ requestSet do
20:          insert cur in cache        /* do cache eviction, if necessary */
21:          for all next ∈ Out-Neighbor(cur) do
22:             if next ∉ resultSet then
23:                resultSet ← resultSet ∪ {next}
24:                nextSet ← nextSet ∪ {next}
25:             end if
26:          end for
27:       end for
28:    end if
29:    requestSet ← φ, currentSet ← φ
30:    exchange(currentSet, nextSet)
31:    hop = hop +1
32: end while
33: output resultSet
```

C. ALGORITHM FOR h-HOP BFS TRAVERSAL IN gROUTING

We provide gRouting's h-hop BFS query evaluation technique in Algorithm 5. We combine BFS traversal with caching (lines 6 and 20) and batch requests (line 18), as shown in Algorithm 5.